WHAT IS CLAIMED IS:

1. A system for selecting open shortest path first (OSPF) aggregates, comprising:

a database for containing data pertaining to candidate OSPF aggregates and corresponding weights; and

an aggregate selector, associated with said database, that selects at least a subset of said OSPF aggregates such that said shortest path length between said particular source and destination subnets resulting from advertisement of a set of weighted aggregates approaches said shortest path length between said particular source and destination subnets irrespective of said advertisement.

2. The system as recited in Claim 1 wherein said aggregate selector treats errors in said shortest path length as having unequal degrees of importance.

3. A method of selecting open shortest path first (OSPF) aggregates, comprising:

storing data pertaining to candidate OSPF aggregates and corresponding weights; and

selecting at least a subset of said OSPF aggregates such that said shortest path length between said particular source and destination subnets resulting from advertisement of a set of weighted aggregates approaches said shortest path length between said particular source and destination subnets irrespective of said advertisement.

4. The method as recited in Claim 3 wherein said selecting comprises treating errors in said shortest path length as having unequal degrees of importance.

5. An autonomous network domain, comprising:

a plurality of routers and interconnecting segments that cooperate to form subnets and paths therebetween; and

a system for selecting open shortest path first (OSPF) aggregates, including:

a database for containing data pertaining to candidate OSPF aggregates and corresponding weights, and

an aggregate selector, associated with said database, that selects at least a subset of said OSPF aggregates such that said shortest path length between said particular source and destination subnets resulting from advertisement of a set of weighted aggregates approaches said shortest path length between said particular source and destination subnets irrespective of said advertisement.

6. The domain as recited in Claim 5 wherein said aggregate selector treats errors in said shortest path length as having unequal degrees of importance.

7. A system for selecting open shortest path first (OSPF) aggregate weights for a particular area, comprising:

a database for containing data pertaining to candidate OSPF aggregates; and

a weight assigner, associated with said database, that assigns, for said OSPF aggregates, weights based on an average distance of subnets in said area for a particular area border router (ABR) of said area.

8. The system as recited in Claim 7 wherein said weight assigner employs a search heuristic to assign said weights.

9. The system as recited in Claim 7 wherein said weight assigner treats errors in path lengths in said area as having unequal degrees of importance.

10. A system for selecting open shortest path first (OSPF) aggregate weights for a particular area, comprising:

a database for containing data pertaining to candidate OSPF aggregates; and

a weight assigner, associated with said database, that employs a search heuristic to assign weights for said OSPF aggregates.

11. The system as recited in Claim 10 wherein said weight assigner treats errors in path lengths in said area as having unequal degrees of importance.

12. A method of selecting open shortest path first (OSPF) aggregate weights for a particular area, comprising:

storing data pertaining to candidate OSPF aggregates in a database; and

assigning, for said OSPF aggregates, weights based on an average distance of subnets in said area for a particular area border router (ABR) of said area.

13. The method as recited in Claim 12 wherein said assigning comprises employing a search heuristic.

14. The method as recited in Claim 12 wherein said assigning comprises treating errors in path lengths in said area as having unequal degrees of importance.

15. A method of selecting open shortest path first (OSPF) aggregate weights for a particular area, comprising:

storing data pertaining to candidate OSPF aggregates in a database; and

employing a search heuristic to assign weights for said OSPF aggregates.

16. The method as recited in Claim 15 wherein said employing comprises treating errors in path lengths in said area as having unequal degrees of importance.

17. An autonomous network domain, comprising:

a plurality of routers and interconnecting segments that cooperate to form subnets and paths therebetween; and

a system for selecting open shortest path first (OSPF) aggregate weights for a particular area in said domain, including:

a database for containing data pertaining to candidate OSPF aggregates, and

a weight assigner, associated with said database, that assigns, for said OSPF aggregates, weights based on an average distance of said subnets in said domain for a particular area border router (ABR) of said area.

18. The domain as recited in Claim 17 wherein said weight assigner employs a search heuristic to assign said weights.

19. The domain as recited in Claim 17 wherein said weight assigner treats errors in path lengths in said area as having unequal degrees of importance.

20. An autonomous network domain, comprising:

a plurality of routers and interconnecting segments that cooperate to form subnets and paths therebetween; and

a system for selecting open shortest path first (OSPF) aggregate weights for a particular area in said domain, including:

a database for containing data pertaining to candidate OSPF aggregates, and

a weight assigner, associated with said database, that employs a search heuristic to assign weights for said OSPF aggregates.

21. The domain as recited in Claim 20 wherein said weight assigner treats errors in path lengths in said area as having unequal degrees of importance.